Massively parallel simulations of Laser-Plasma Acceleration for e-/e+ Collider Concepts

presented by David L. Bruhwiler[1]

COMPASS team & close collaborators:

J.R. Cary,[1,2] E. Cormier-Michel,[3] B. Cowan,[1] V. Decyk,[4] E. Esarey,[3] C.G.R. Geddes,[3] C. Huang,[4] W.P. Leemans,[3] W. Mori,[4] K. Paul,[1] C.B. Schroeder,[3] F. Tsung[4] J.-L. Vay[3]

VORPAL optimization and scaling:

K. Amyx,[1] T. Austin,[1] G.I. Bell,[1] P. Messmer,[1] P. Mullowney[1] & B. Norris[5]

VACET collaborators:

W. Bethel,[3] J. Jacobsen,[3] Prabhat,[3] O. Rubel,[3] D. Ushizima[3] & G. Weber[3]

ComPASS SciDAC Review; Rockville MD; April 21, 2009

1. Tech-X Corporation
2. University of Colorado
3. Lawrence Berkeley Lab
4. UCLA
5. Argonne National Lab

Primary support: DOE Office of Science, Office of High Energy Physics, SciDAC-2
Additional support: DOE SBIR program (HEP, ASCR); NNSA / NA-22; DARPA
DOD SBIR program (AFOSR, OSD); Tech-X Corp. customers

Office of Science

SciDAC Scientific Discovery through Advanced Computing

BELLA* Project Underway: World-Leading Facility for Laser-Based Accelerator Science

- High rep rate (1 Hz), Petawatt class laser (>40 J in < 40 fs)
- 10 GeV beam in 1 meter
  - Module for future collider
  - Positron source, femtosecond x-ray FEL

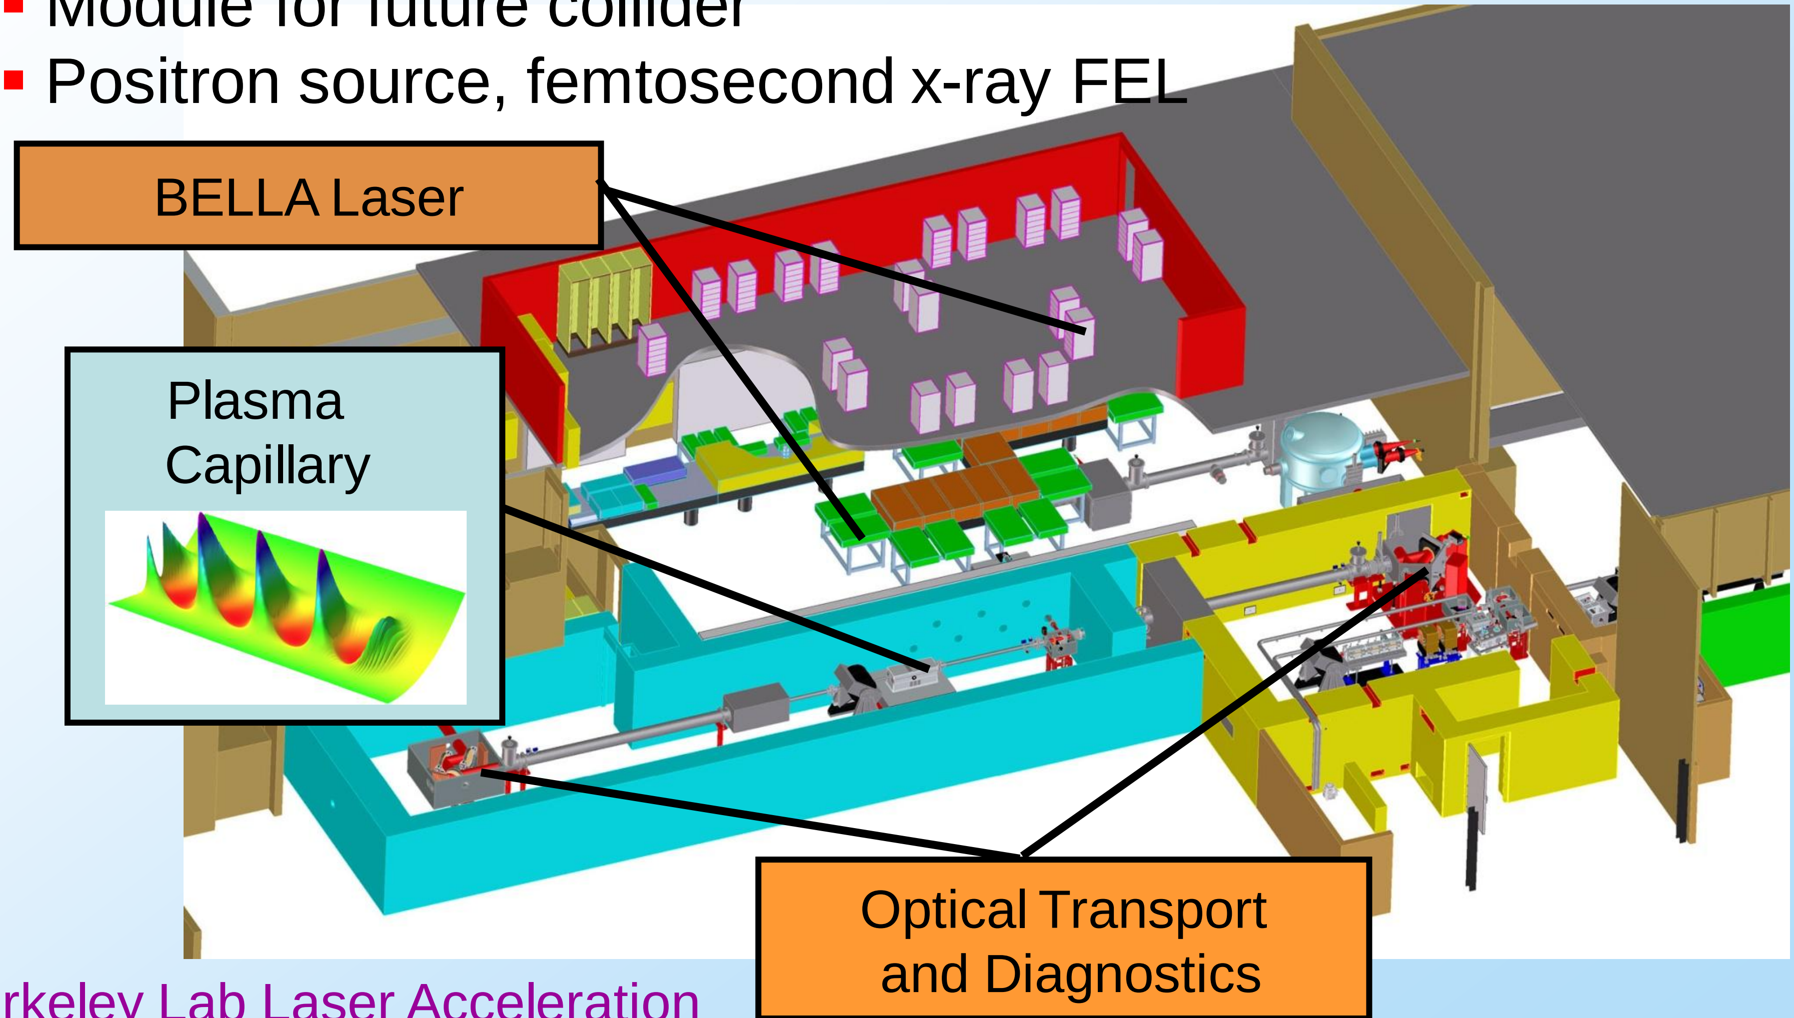

*Berkeley Lab Laser Acceleration

Regimes for a 300J laser

	Self-guiding		External-guiding	
	Self Injection I*	Self Injection II**	Self Injection**	External Injection**
Laser				
a0	43	5.8	3.5	2
Spot [μm]	9	50	70	101
Duration [fs]	30	110	155	224
Plasma				
Density [cm^{-3}]	1.5×10^{19}	2.7×10^{17}	8.2×10^{16}	2.2×10^{16}
Length [cm]	0.25	22	100	500
e- Bunch				
Energy [GeV]	**4**	**13**	**25**	**53**
Charge [nC]	14	2	1.8	1.5

Maximum electron energy

Strongly nonlinear *Full PIC Laboratory*	**Nonlinear** Full PIC Boosted & QuickPIC Lab.		**Weakly Nonlinear** Full PIC Boosted & QuickPIC Lab.

* Gordienko and A.Pukhov, Phys Plasmas B, 043109 (2005).
** W. Lu et al., Phys. Rev. ST Accel. Beams **10**, 061301 (2007).

Scaling with density used to simulate m-scale 10 GeV stages for BELLA

- Order(GHours) for direct sim.
- Scaling allows shorter runs:
  - Constant: L_{laser}/λ_p, w_0/λ_p, a_0
  - Vary density: energy~$1/n_e$

Wake scales with density

Scaled simulations at a=1

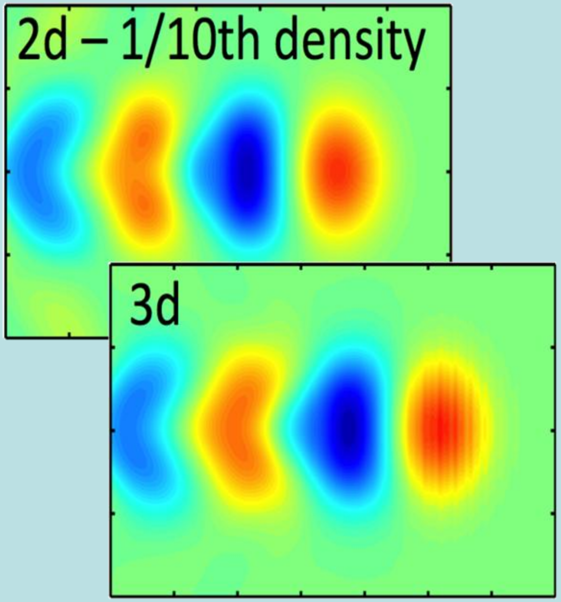

- Beam loading scaling allows prediction of stage charge

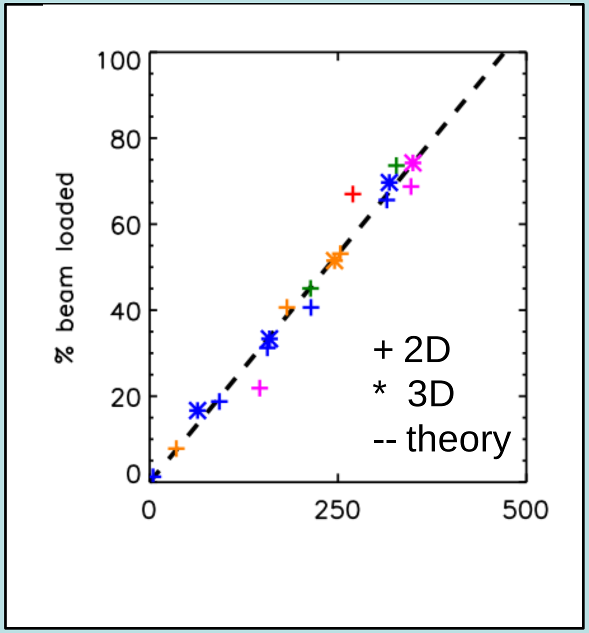

density & k_pL:	$k_p\sigma_r$ = 0.5	1	1.8
k_pL =2, a_0=1 $n_0 = 10^{18}$ cm^{-3}			+
k_pL =2, a_0=1 $n_0 = 10^{19}$ cm^{-3}	+*	+*	+*
k_pL =1, a_0=1.4 $n_0 = 10^{19}$ cm^{-3}	+		

- Particle bunch shaping and plasma taper increase gain, reduce energy spread
- Quasi linear regime allows symmetric acceleration of positron buch

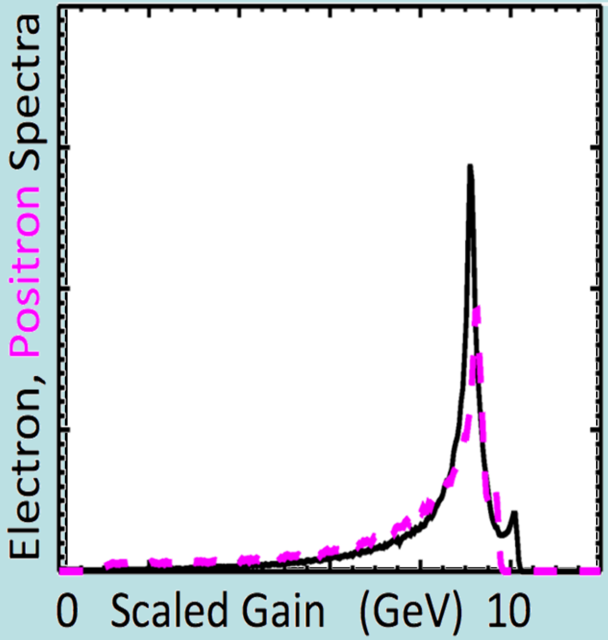

9 GeV energy gain

4% energy spread (FWHM)

0.7 m acceleration length

E. Cormier-Michel et al. 13th AAC workshop proc. (2008)

Simulations tune laser and bunch: Efficient stage near $k_pL = 1$

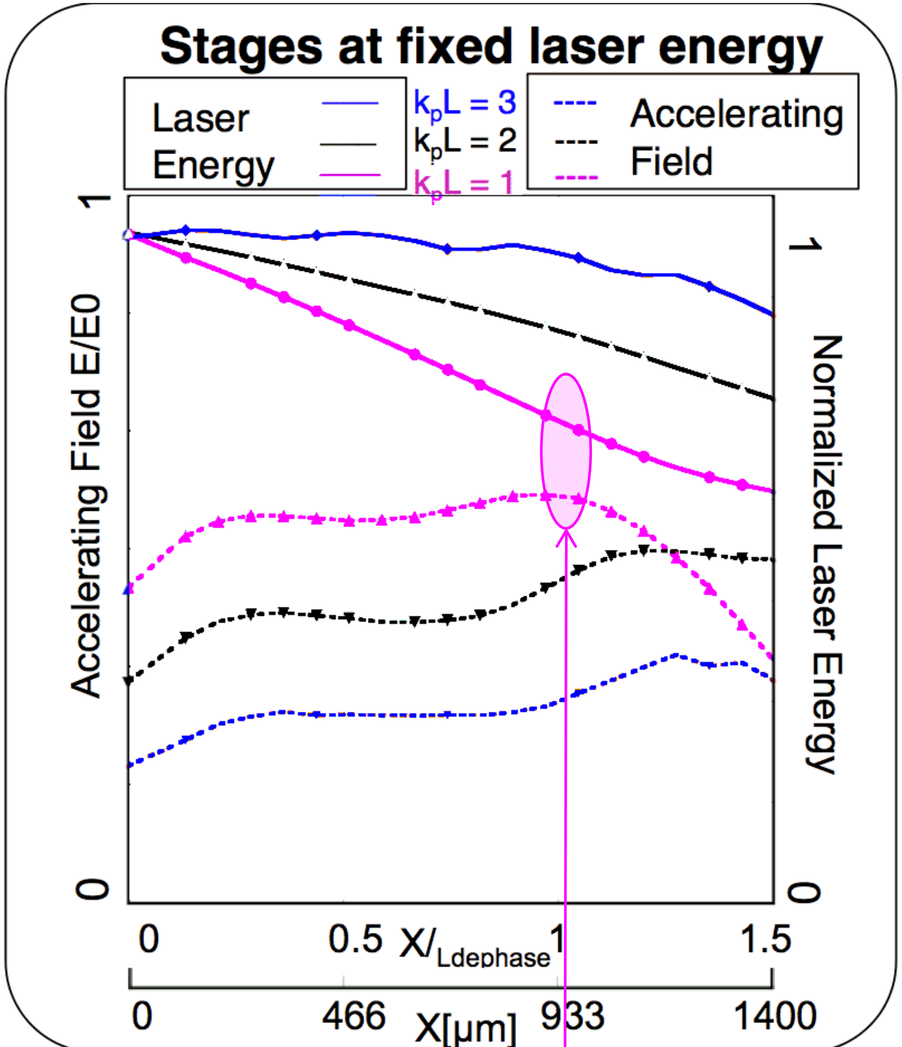

At $k_pL = 1$: 16 GV/m for 10 GeV, efficient depletion at dephasing

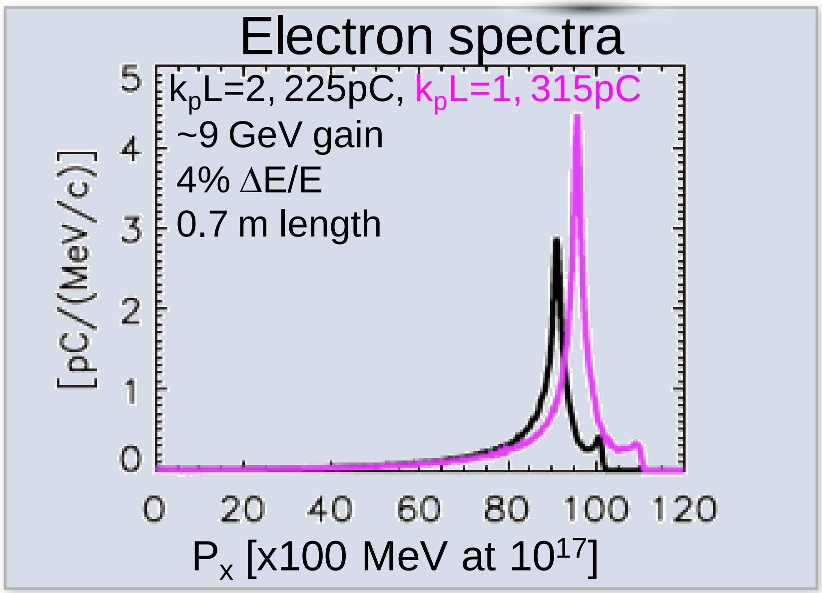

Resolve laser depletion, broadening

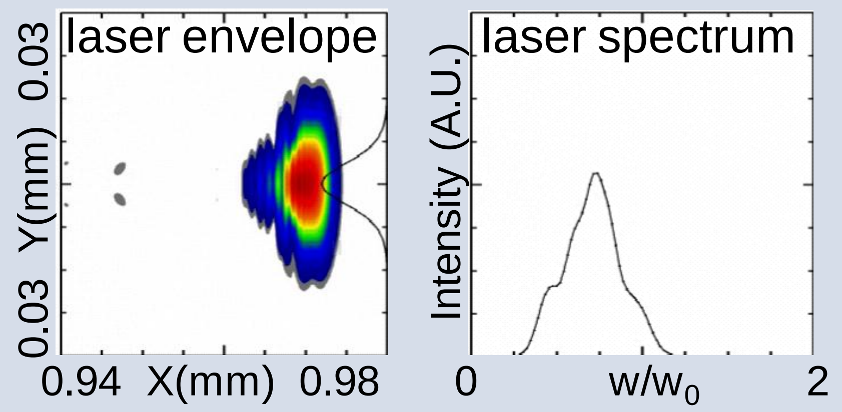

- Further increase efficiency – use shaped bunches and laser pulses

Boosted-frame concept; 1D proof-of-principle in VORPAL

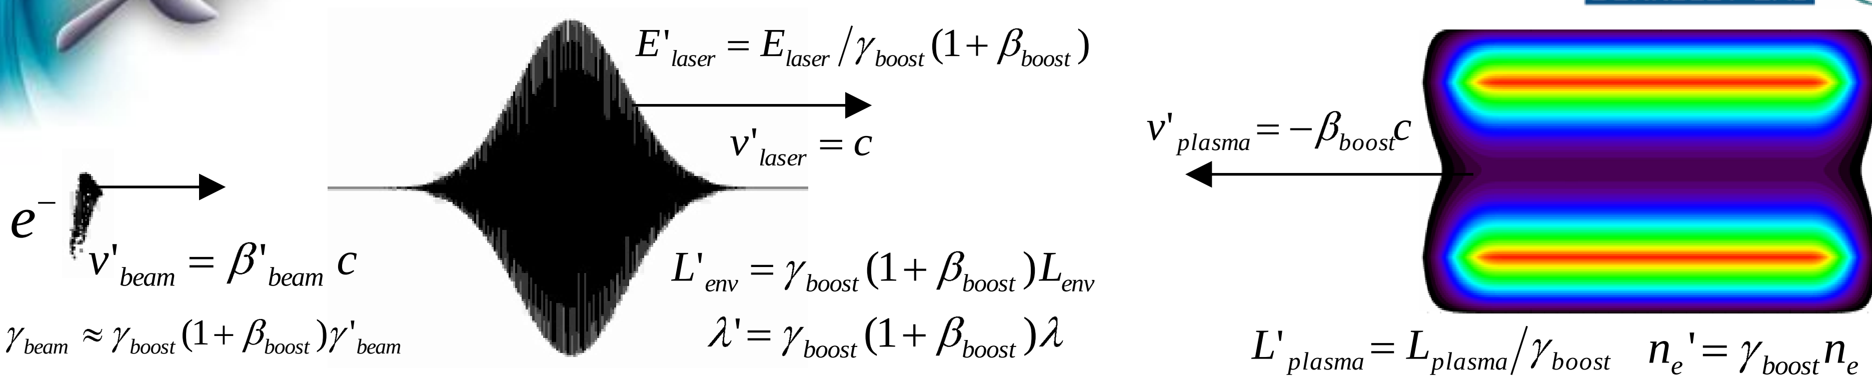

- Could enable new simulation regimes; dramatically speed up existing simulations
- Grid size & resolution are equivalent to standard lab frame runs
- 2D example shown below: $n_e = 6 \times 10^{16}$ cm^{-3}; $L_{deph} \sim 2.4$ m; $a_0 = 1$; $E_{peak} \sim 11$ GeV (lab)
  - agrees with scaling estimate
- Speed-up of 2,000x
- Total-field / scattered-field emitter added to VORPAL
- More work required
  - improved noise reduction
  - automated set-up, diagnostics
  - validation and testing

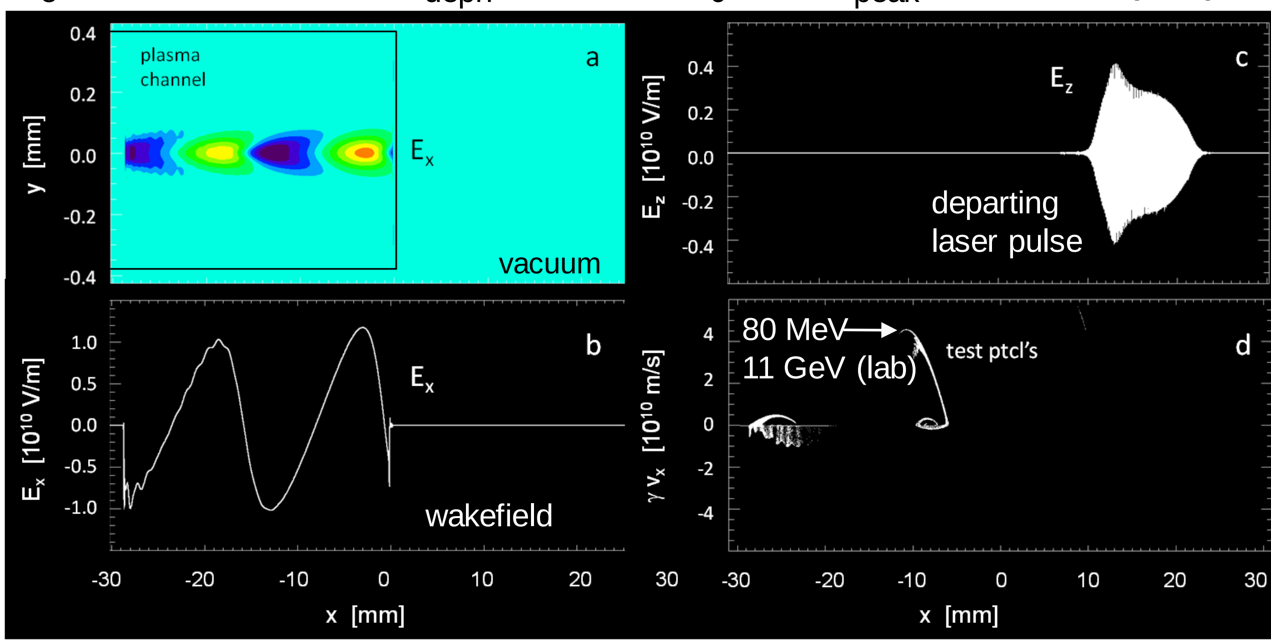

D. Bruhwiler *et al.*, Proc AAC 2008

Primary support, DOE/HEP SBIR program.

Ultra-fast 3D simulations in the Boosted Frame : 20 to 500x speedups

Supporting experimentalist with current experiments...

Lawrence Livermore National Laboratory

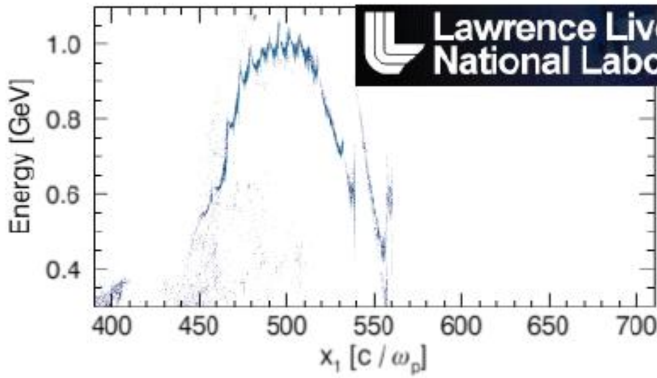

Imperial College London

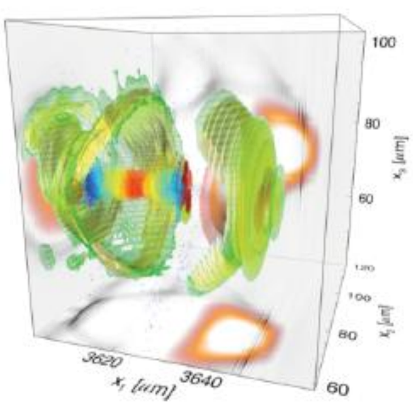

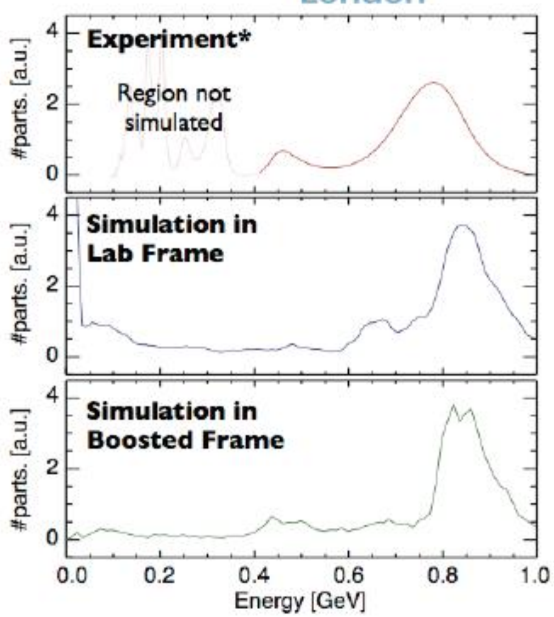

...and designing the next generation of LWFA stages

12 GeV self-injected beam with ~1nC

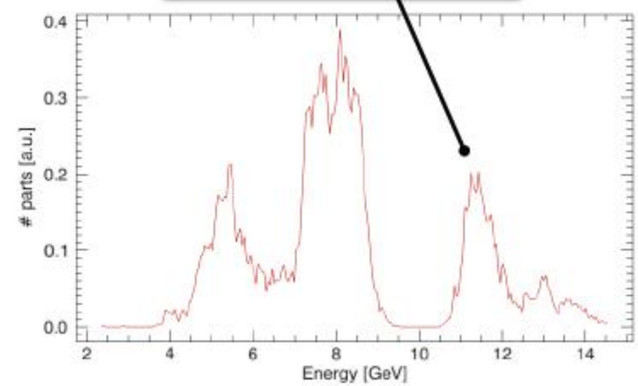

2D boosted-frame LWFA simulations now working successfully in WARP

- New electromagnetic solver implemented in Warp (SBIR funding)
  - scaling test (3-D decomp)

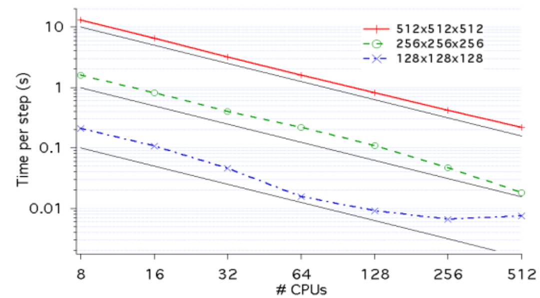

# procs decomposition	# cell, particles	Efficiency
256 (8×8×4)	$1{,}024^2 \times 512$ 100M	1.
512 (8×8×8)	$1{,}024^3$ 200M	1.04
1024 (8×8×16)	$1{,}024^2 \times 2{,}048$ 400M	1.12

- Applied to modeling of one stage of LWFA (2-D for now, 3-D to follow)

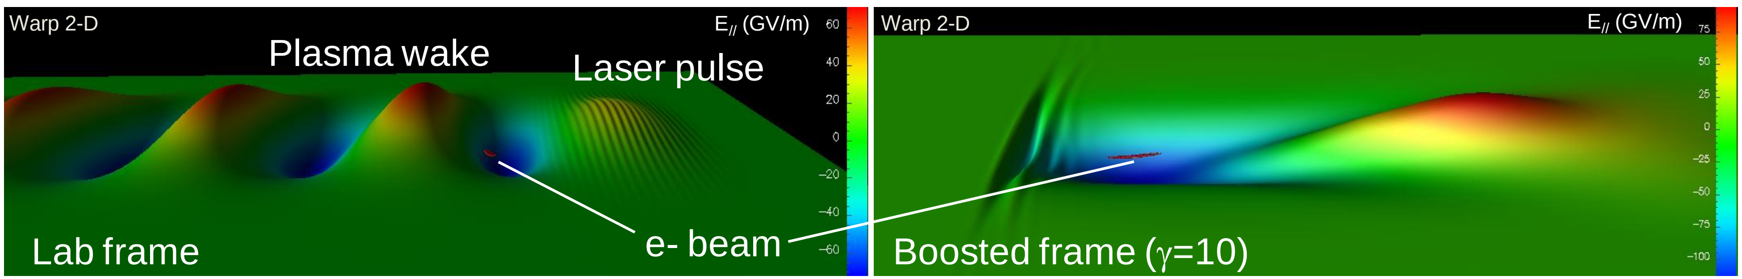

collaboration with LBNL's LOASIS group (lead by Wim Leemans)

Average beam energy and CPU time vs position in lab frame

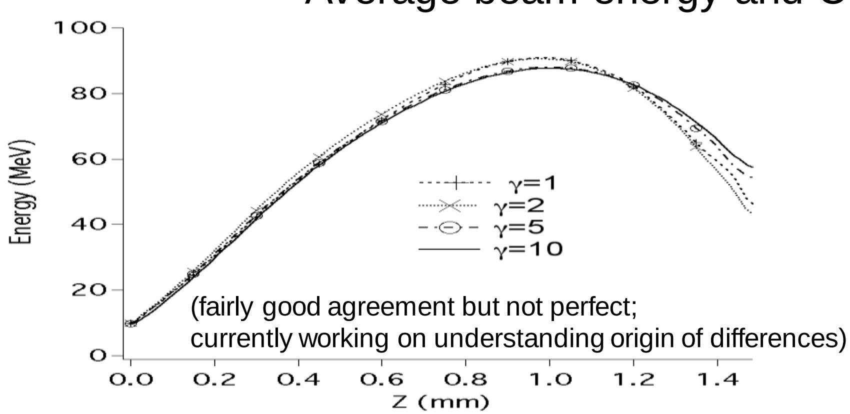

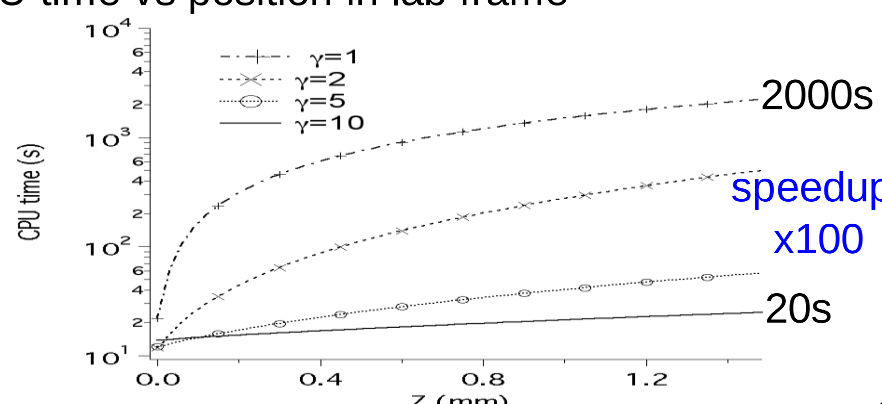

Envelope model simulates 10 GeV at scale; 3D downramp injection of plasma electrons

- **Don't resolve laser wavelength or frequency**
- **Particles respond to ponderomotive force**
- **Full plasma dynamics via standard PIC**
  - high-order particle shapes
  - current smoothing
  - absorbing boundary conditions, etc.

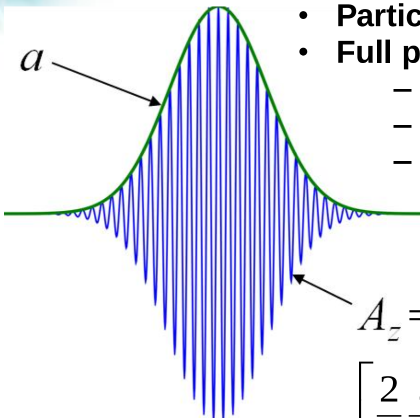

$$A_z = \mathrm{Re}\left[a e^{i(\omega t - k_0 x)} \right] = \mathrm{Re}\left[a e^{-ik_0 \xi} \right]$$

$$\left[\frac{2}{c}\frac{\partial}{\partial \tau}\left(\frac{\partial}{\partial \xi} - ik_0 \right) + \nabla_\perp^2 \right] a = -\mu_0 \chi a$$

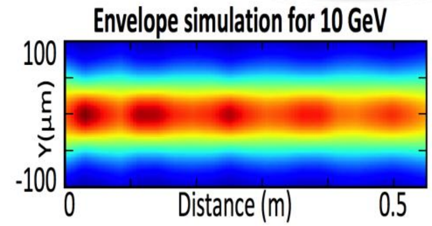

correct betatron period

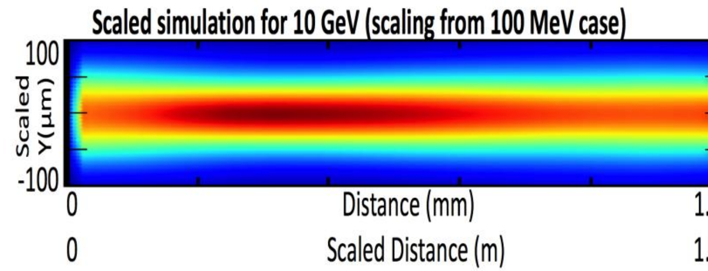

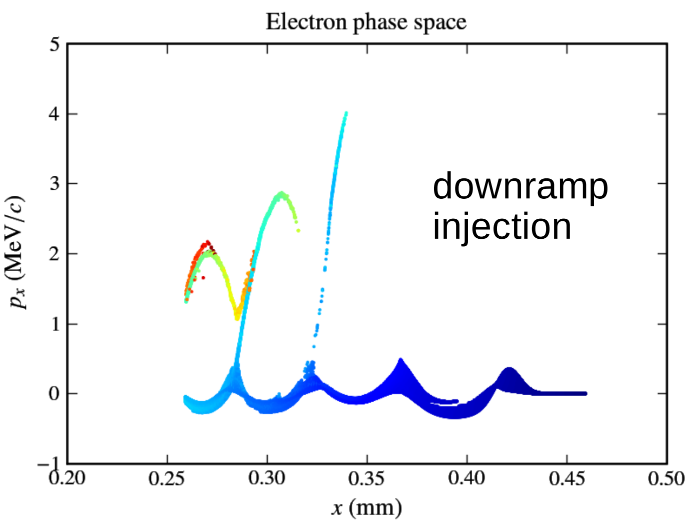

- Challenges for envelope simulations of Bella stages
  - spectral broadening due to pump depletion limits distance
  - small, low-emittance injected bunch requires fine mesh
- Speed-up factor $\approx .1^*(\lambda_p/\lambda_0)^2 \sim 1/n_e$ (typically 10x to 100x)
  - transverse resolution same as for standard PIC
  - Trilinos for implicit solve; 512 Franklin cores is typical
- Ideal algorithm for simulating downramp injection
  - laser spot at back of gas jet → propagation of converging pulse over $>Z_R$ → wide transverse simulation domain
  - full PIC simulations effectively limited to 2D
  - 2D envelope simulations agree with full PIC → 3D next

Mora & Antonsen, Phys. Plasmas 1997; Gordon *et al.*, IEEE Trans. 2000
Messmer & Bruhwiler, PRSTAB 2006; B. Cowan *et al.*, Proc AAC 2008

VORPAL laser-envelope model successfully benchmarked with time-explicit PIC

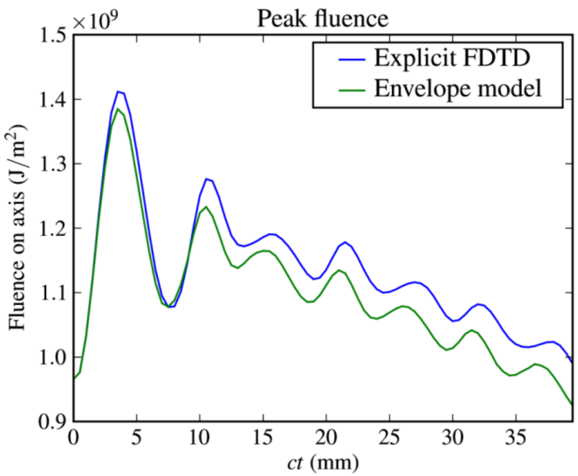

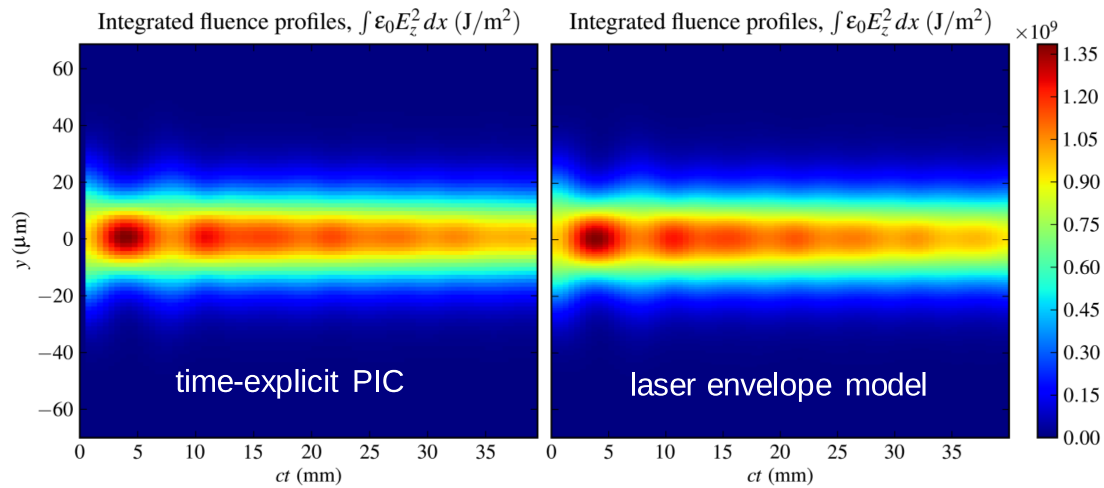

- Good agreement; extended channel propagation, betatron oscillations; pump depletion
- 2D, scaled 10 GeV parameters; $n_0 = 10^{24}\ m^{-3}$; $a_0 = 1$; speed-up of 18x

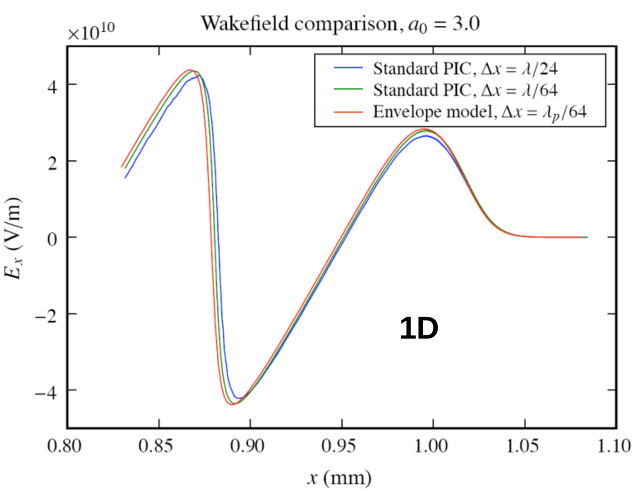

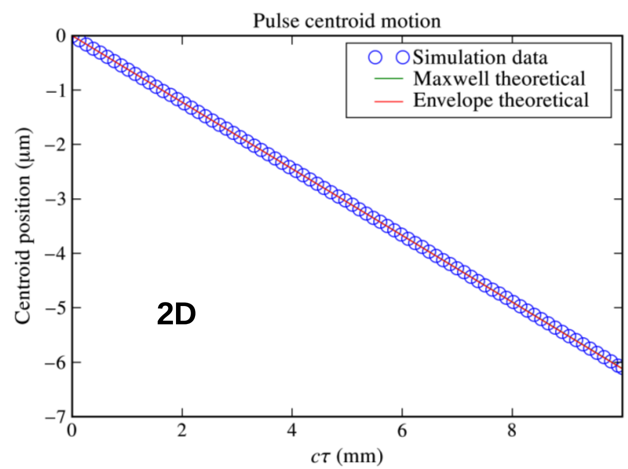

- **Converged plasma wakefield**
  - $\lambda_p/32$ for envelope model
  - $\lambda_0/64$ for time-explicit PIC
- **Correct group velocity**
  - no Yee dispersion errors

B. Cowan *et al.*, Proc AAC 2008

TECH-X CORPORATION

QuickPIC: +10 GeV in self-guided regime

Laser spot evolution

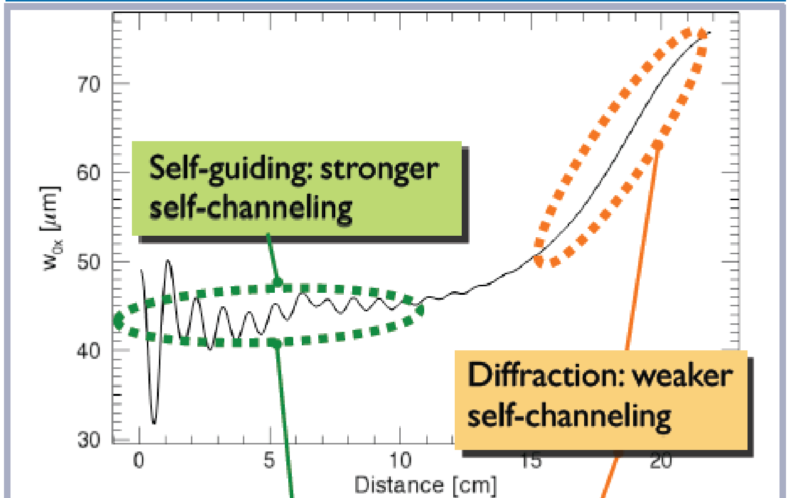

Main results

Two regimes for laser propagation:

- Self-guiding propagation regime until 10 cm
- Depletion leads to diffraction after 10 cm

Accelerating gradient in good agreement with theory

- QuickPIC: 0.8 GeV/cm
- Theoretical: 0.6 GeV/cm

Spectral evolution

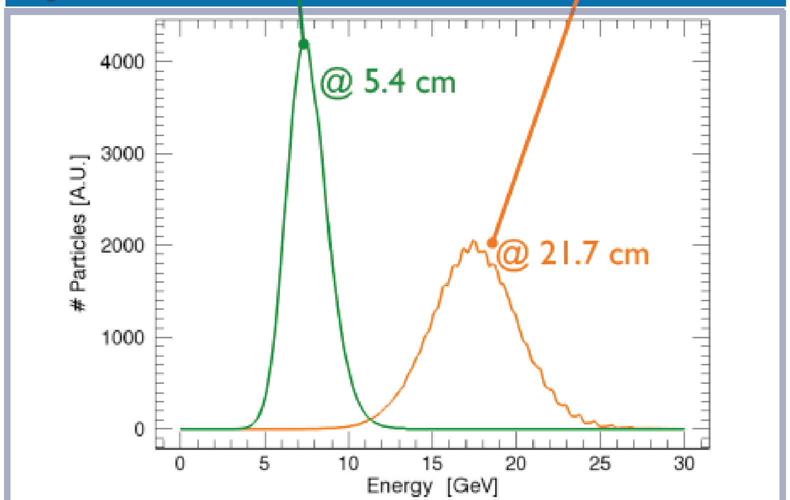

Phase-space evolution

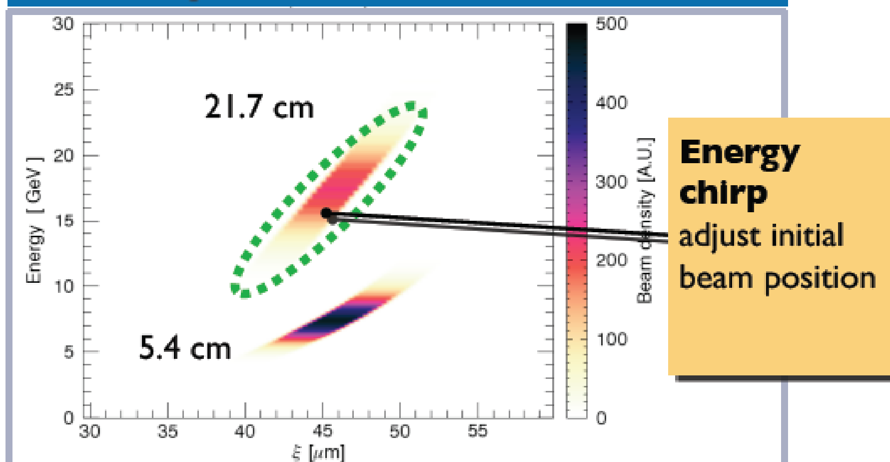

QuickPIC:weakly nonlinear regime scales well out to 100 GeV

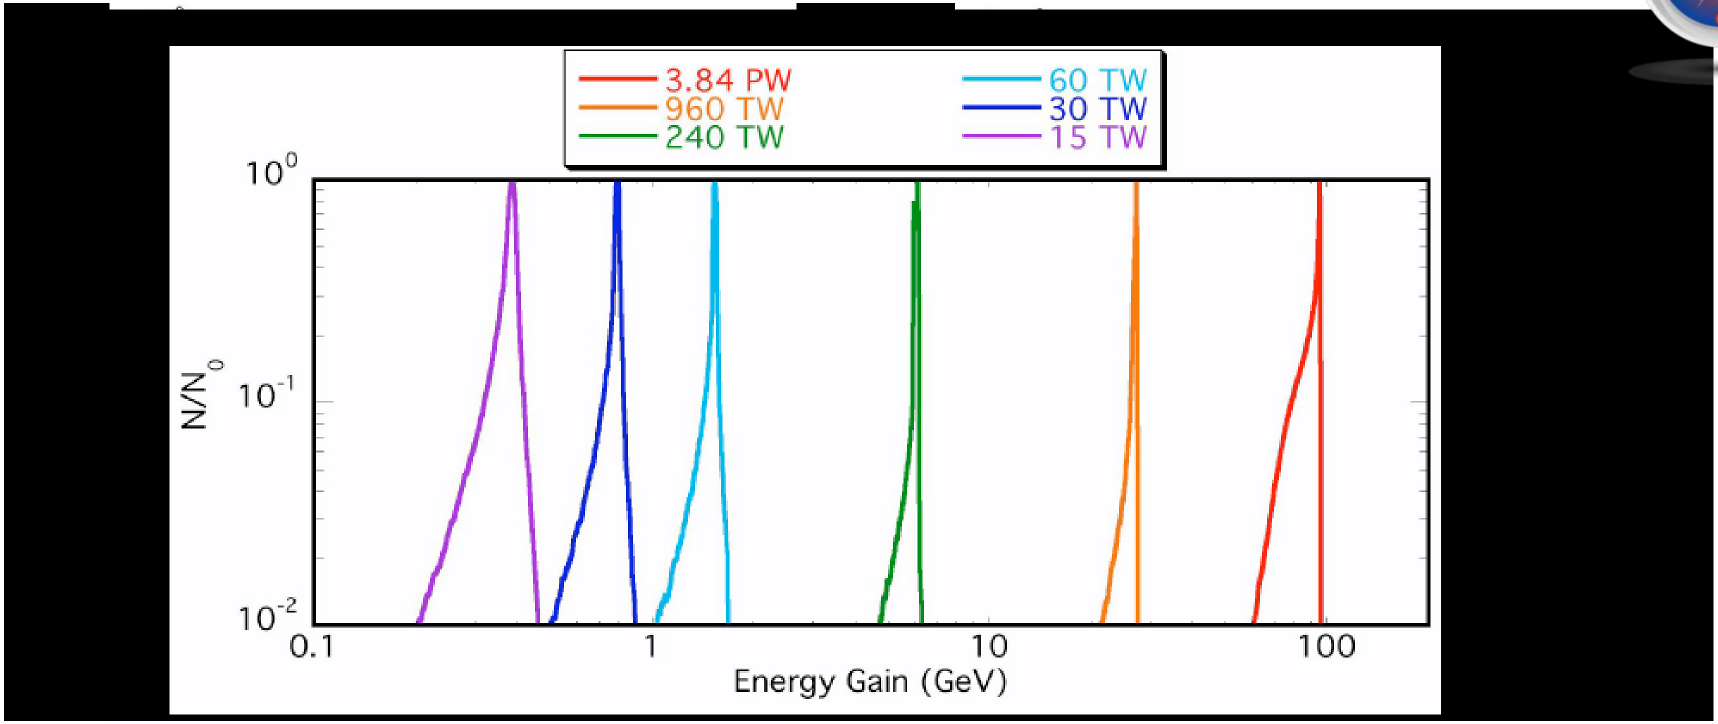

- N_0 increases with output energy.
- The emittance remains relatively constant throughout all of the simulations.
- To reduce the energy spread an exact theory is required for beam loading and for the evolution of the laser after hundreds of Z_R.

Self trapped experiments: Physics of percent energy spread, verify scaling

- Laser channeling: first low $\Delta E/E$ beams
  - 10 TW laser, 2mm plasma @ 2×10^{19}/cc

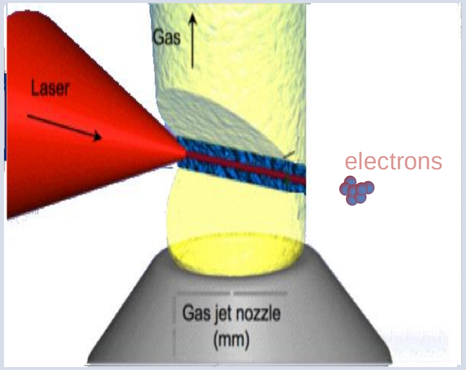

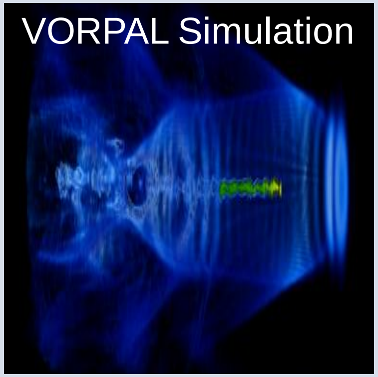

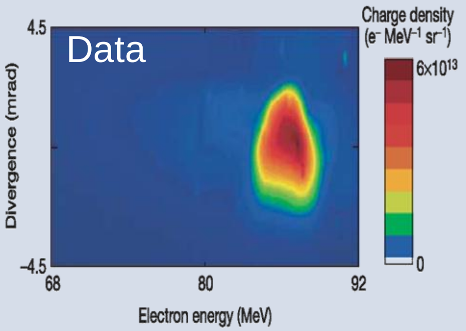

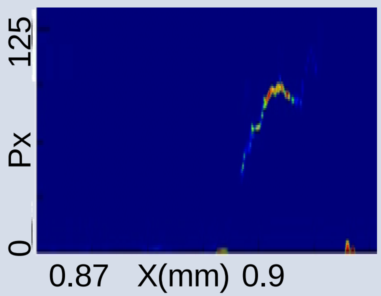

- 3 mrad divergence, $\Delta E/E$ 4%, E_{peak}~170MeV

Geddes et al., Nature 2004; SciDAC Review 2009

- Capillary channels+low n_e=GeV in 3 cm
  - 40 TW laser, 3cm plasma@$4\text{-}5\times10^{18}$/cc

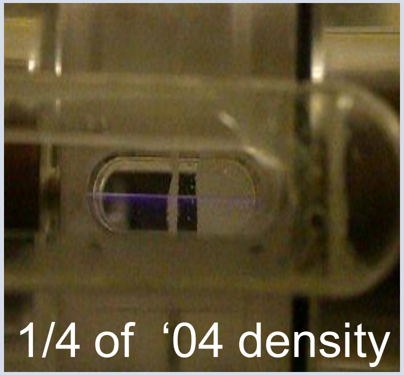

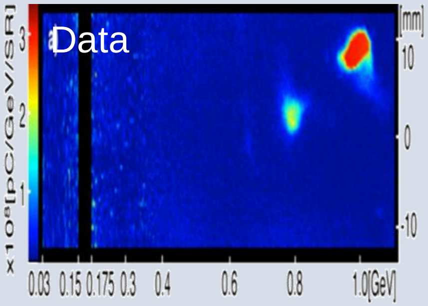

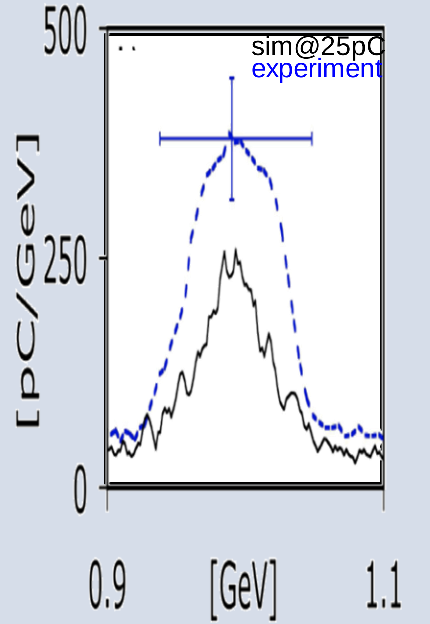

- 1 GeV beams, stable beams at 0.5 GeV

Leemans et al., Nature Phys 2006

- Simulations show physics of self trapping, production of narrow ΔE:
  - beam loading and dephasing – MHour simulations in 3D
- Bunch energy scales as expected with laser, plasma
- 100 MeV 3D production runs at 11k processor/36 hr, 2D 256 processor/1 hr
  - meets near-term goal of providing experimental feedback on a scale of hours

Plasma downramp trapping: all-optical low-ΔE, low-ε injector

Experiments focus on gas jet downramp

Low momentum spread

single shots

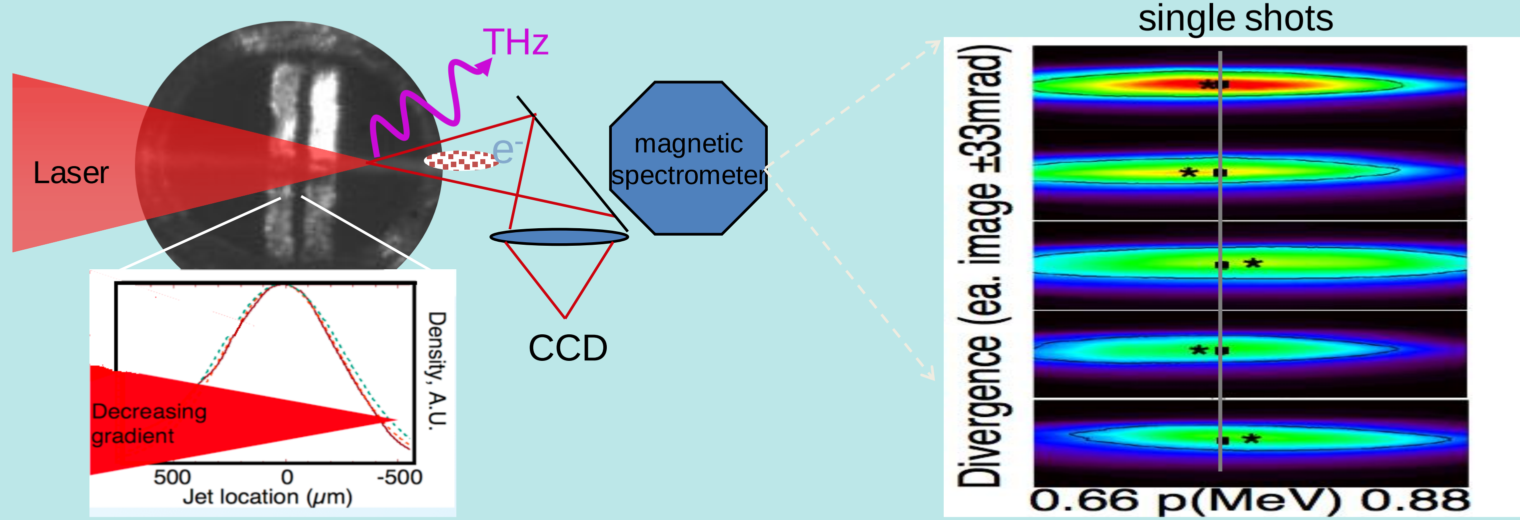

- Validate: VORPAL simulations vs. diagnostics
  - MeV momentum, Δp ~ 200 keV/c
  - 20-50 keV/c transverse momentum
  - 70% laser transmission
  - Ultrashort bunch – THz diagnostics
- Physics: ramp controlled trapping threshold
- Ramp → channel: low ΔE at high E
- Experiments in progress : downramp and also colliding pulse to optimize injector

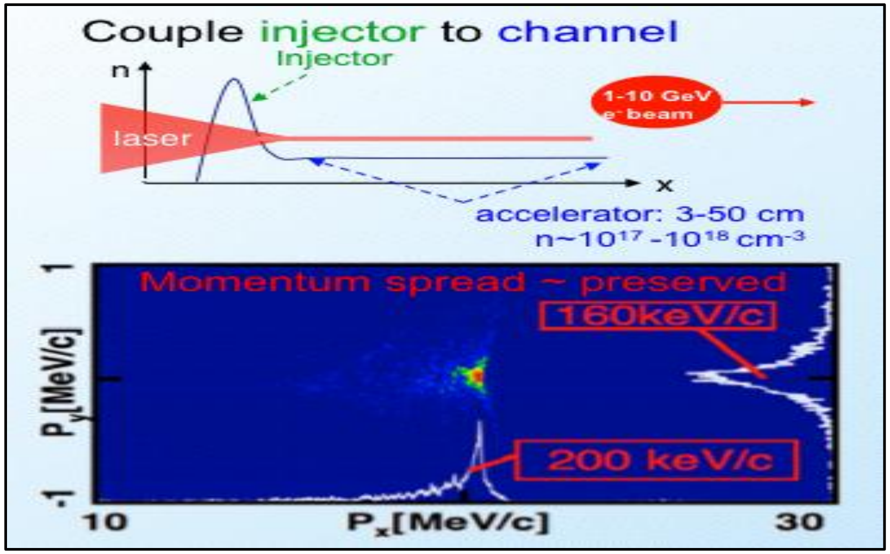

Geddes et al, PRL 2008, JPCS 2008.; G. Plateau, proceedings PAC 2007; Concept: Bulanov PRE 1998

Improving PIC momentum spread accuracy: more accurate modeling of bunch emittance

- Discretization introduces interpolation, error
  - Unphysical temperature, emittance
  - Slow improvement with resolution
- Momentum errors - reduced 100x by:
  - High order spline interpolation
  - Current smoothing
  - Simulated temperature close to expt.
  - Reduces unphysical trapping
- Divergence close to 100 MeV experiment
  - Improves design of low-emittance stages
  - Further work required for collider emittances

Quadratic weight (blk) reduces error

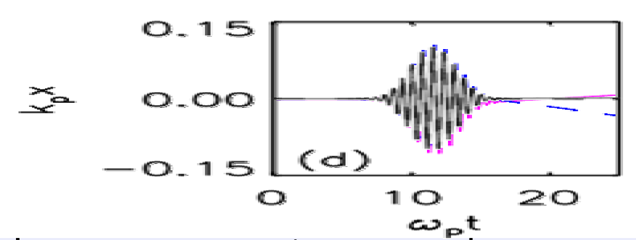

Increase momentum spread accuracy

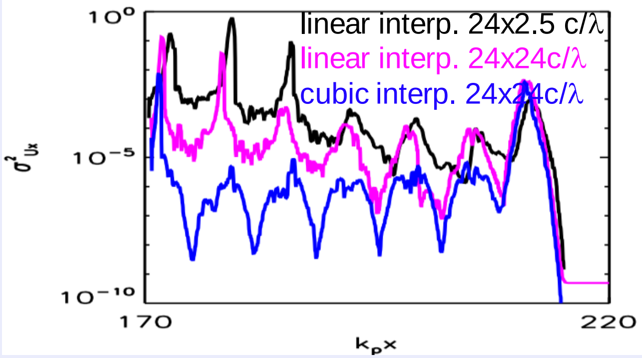

Accurately model bunch divergence

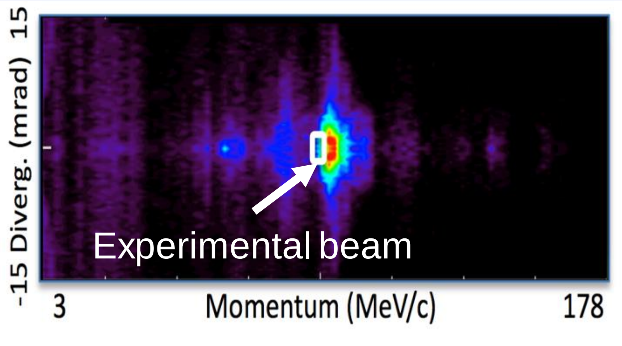

E. Cormier-Michel et al., PRE 2008, Geddes et al. Proc AAC 2008 & SciDAC Review 2009, Bruhwiler et al. in preparation

Laser mode controls transverse field, sets bunch emittance matching

- Emittance matched bunch radius << λ_p for nonlinear & Gaussian-laser linear
  - Laser mode shaping increases matched bunch radius & loading efficiency
  - Fields can be shaped to compensate emittance

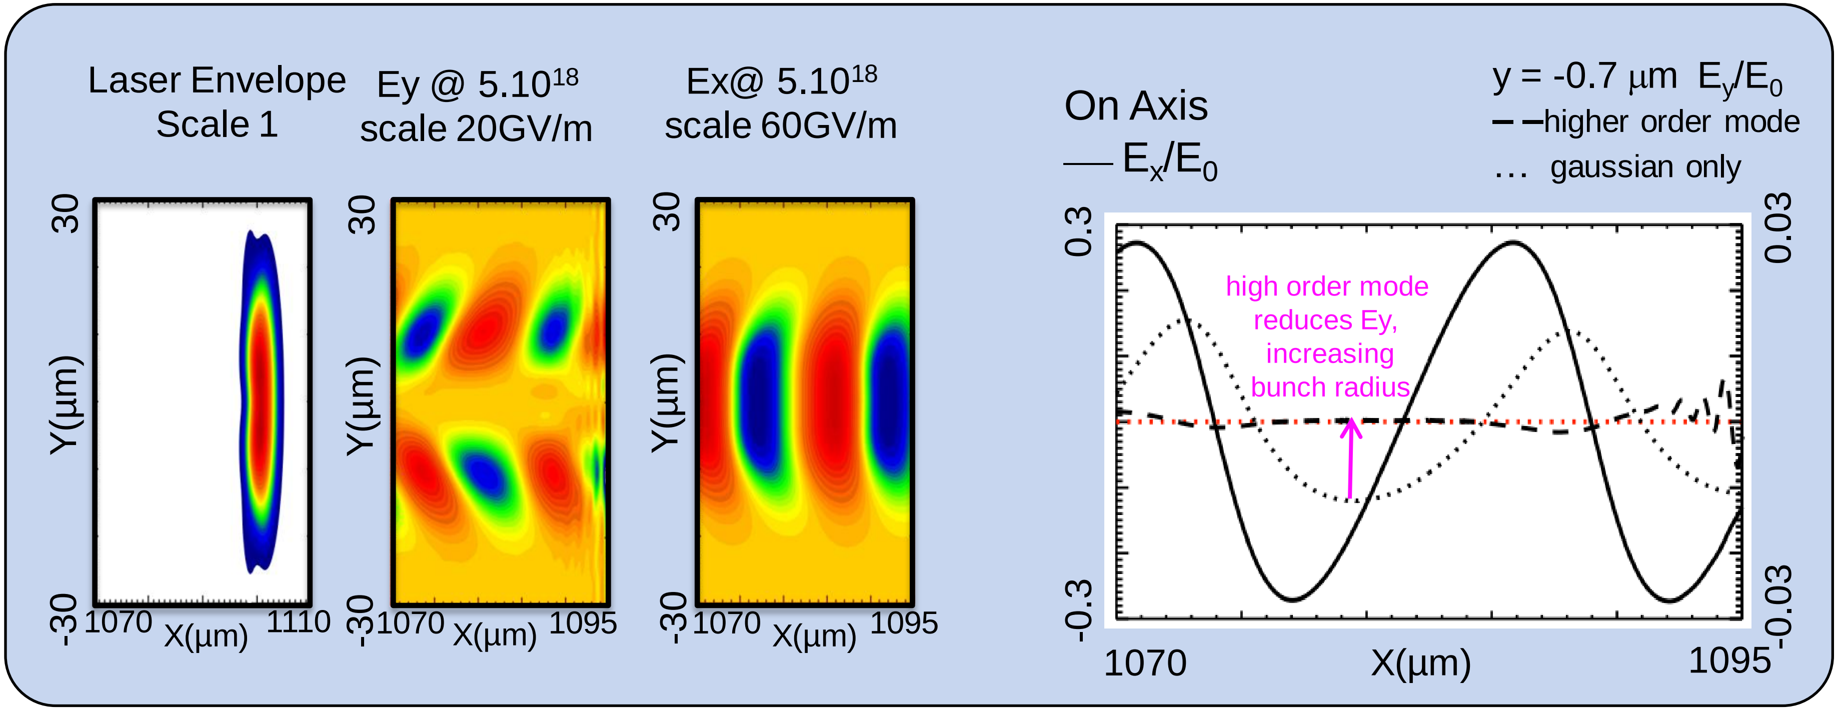

- Detailed emittance modeling requires integration of momentum accuracy and potentially additional models such as mesh refinement, radiation, scattering models

Massively parallel Full PIC results for a 300J laser

Strongly nonlinear (30fs) :: Laboratory

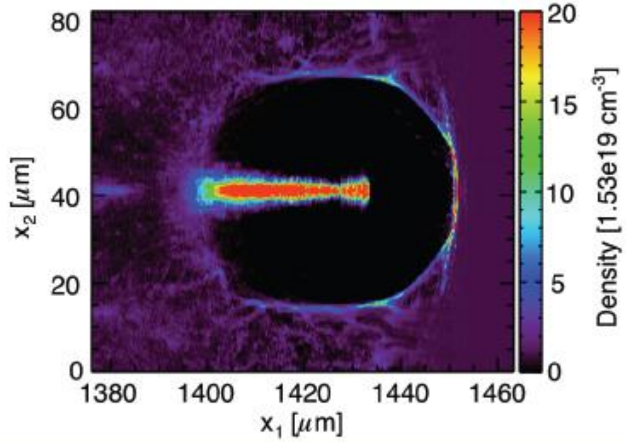

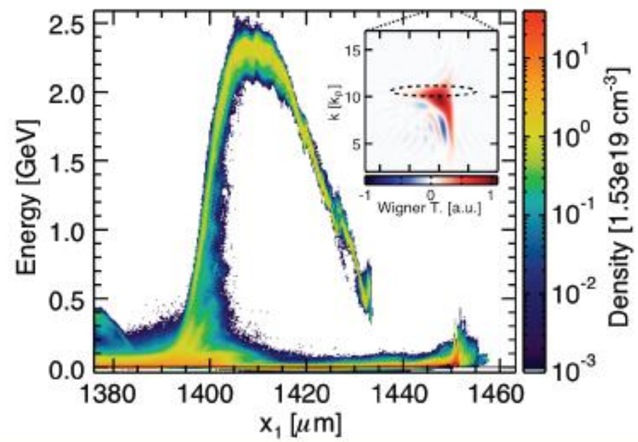

Nonlinear (110fs) :: Boosted

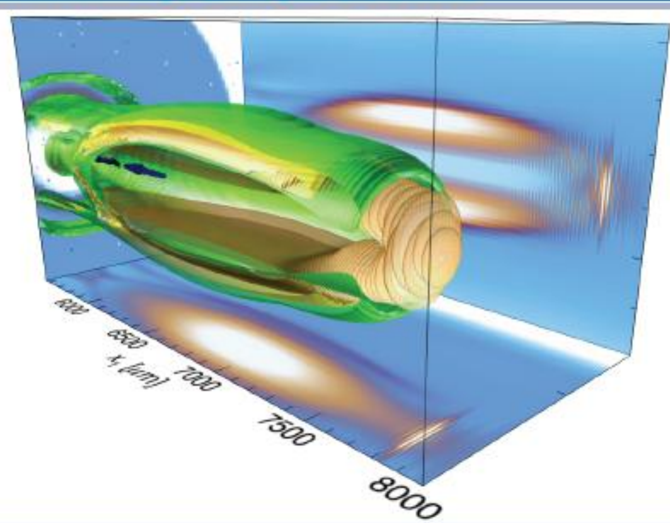

Weakly nonlinear (225fs) :: Boosted

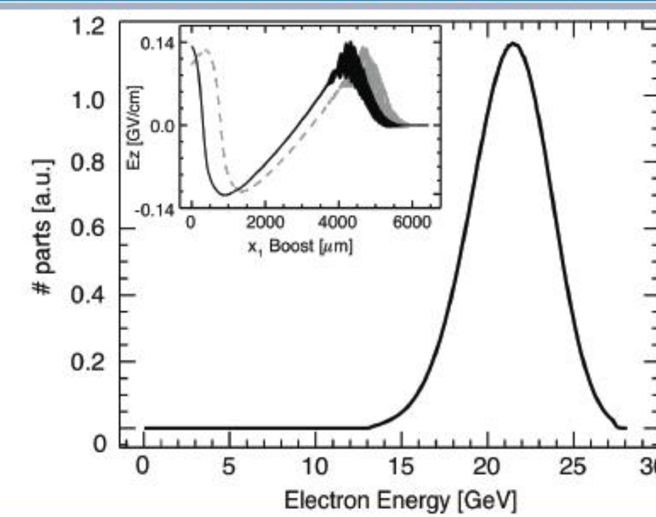

Collaborative development & use of analysis & viz tools; VisIT & FastBit

Fast bit indexing, query, and tracking*

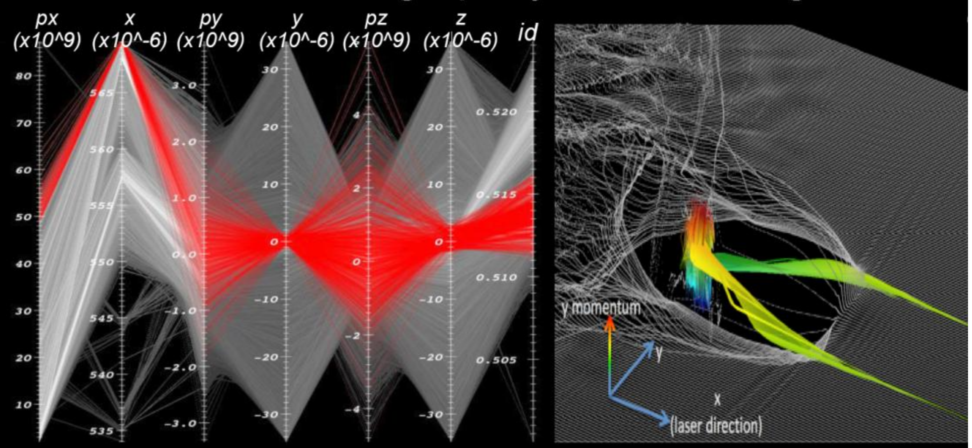

Injection dynamics and interactive exploration of TB datasets

VisIT 3D visualization

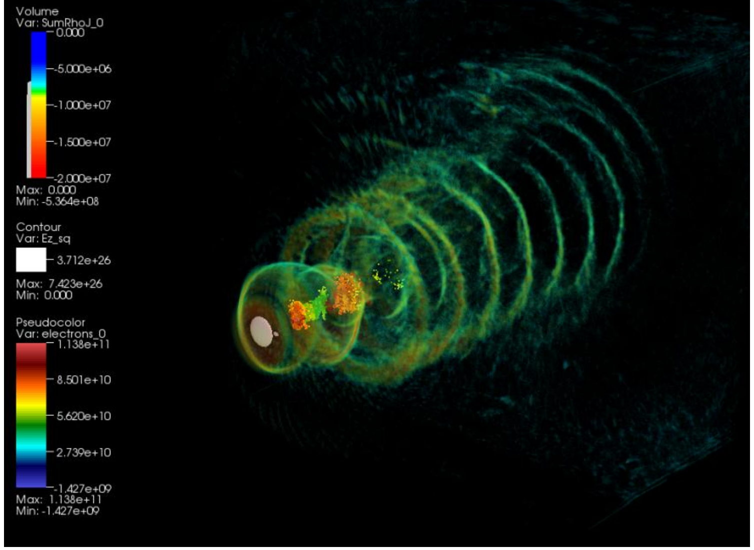

* O. Rubel *et al*., accepted in Supercomputing (2008).

** D. Ushizima *et al*., ICMLA (2008), submitted.

K.J. Wu *et al*., to appear in SciDAC Review (2009).

C.G.R. Geddes *et al*., to appear in SciDAC Review (2009).

Fuzzy clustering in 6D phase space + peak detection**

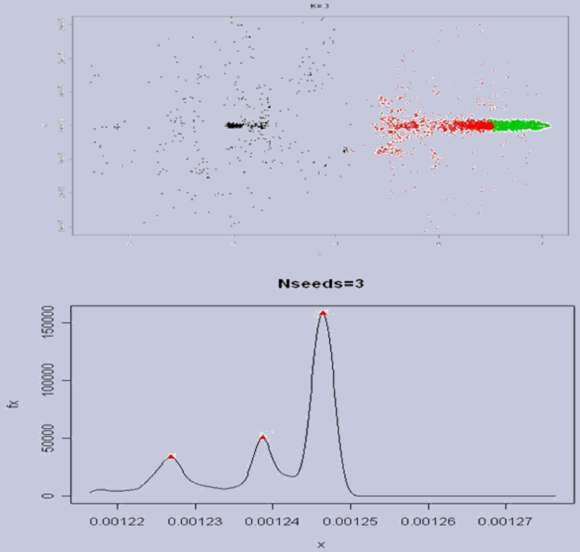

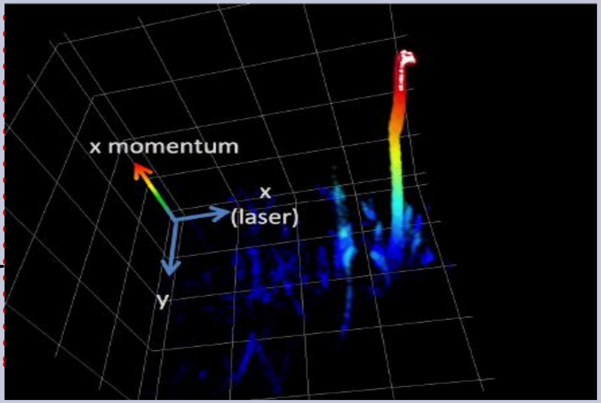

Automated beam detection

Particle tracking in OSIRIS

Relevant physics associated with small subset of particles
Record detailed 7D phase-space of "interesting" particles

- **Technically challenging**
  Subset of ~10^3 particles in ~10^9
- Storing information for every particle not feasible
  - 10^4 iter. × 10^9 part. ⇒ ~ 500 TB

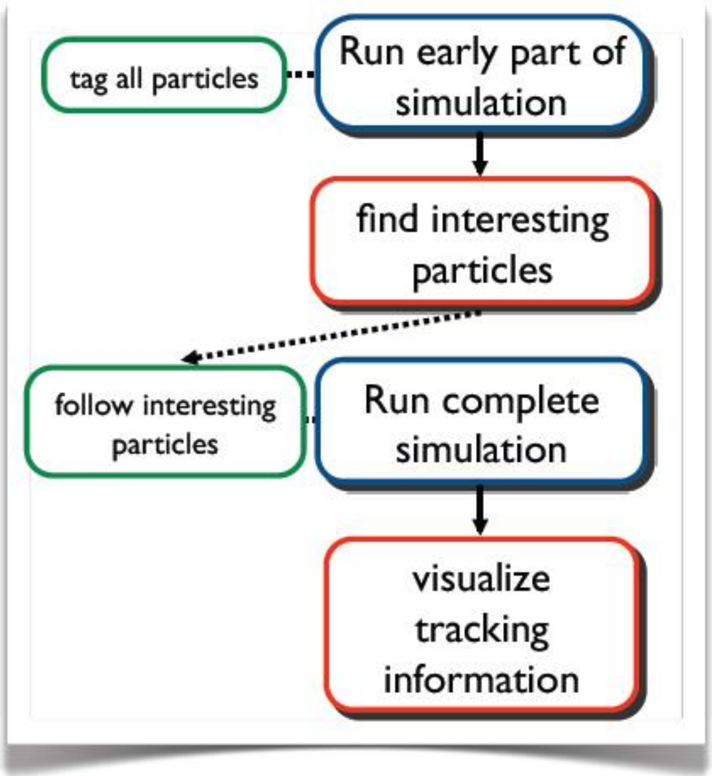

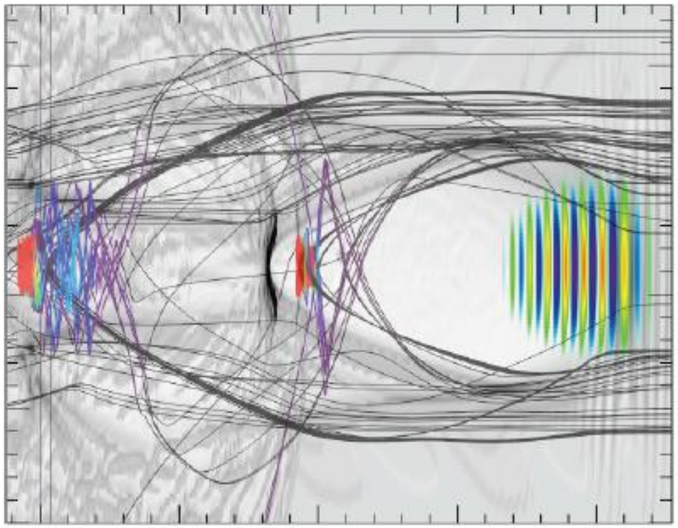

Radiation diagnostics using particle tracking

Post-processing particle tracking

Extract particle tracking

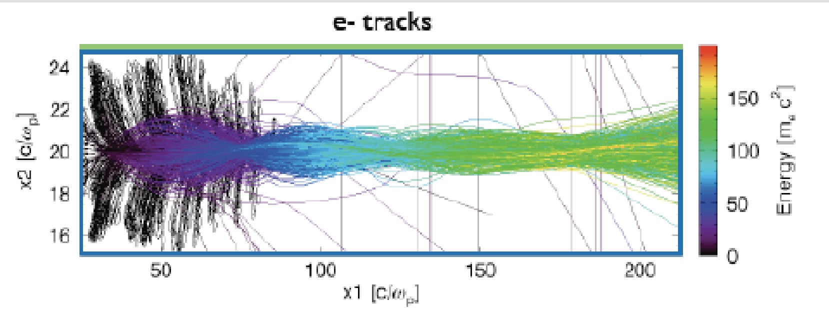

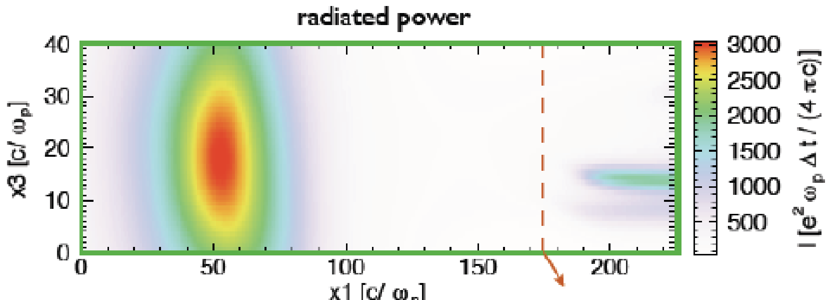

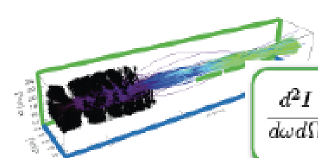

$$\frac{d^2I}{d\omega d\Omega} = \frac{e^2}{4\pi c}\left|\int_{-\infty}^{\infty} \frac{\vec{n}\times[(\vec{n}-\vec{\beta})\times\dot{\vec{\beta}}]}{(1-\vec{n}\cdot\vec{\beta})^2} e^{i\omega(t'-\vec{n}\cdot\vec{r}(t')/c)}dt\right|^2$$

Power and frequency spectra

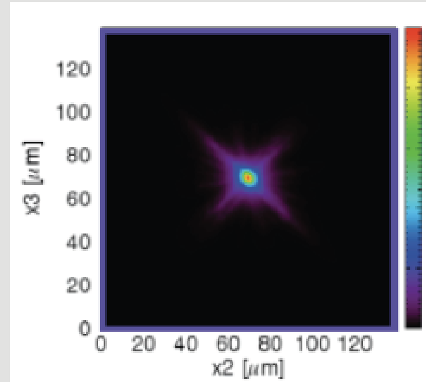

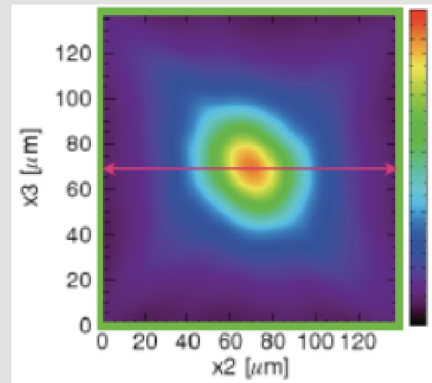

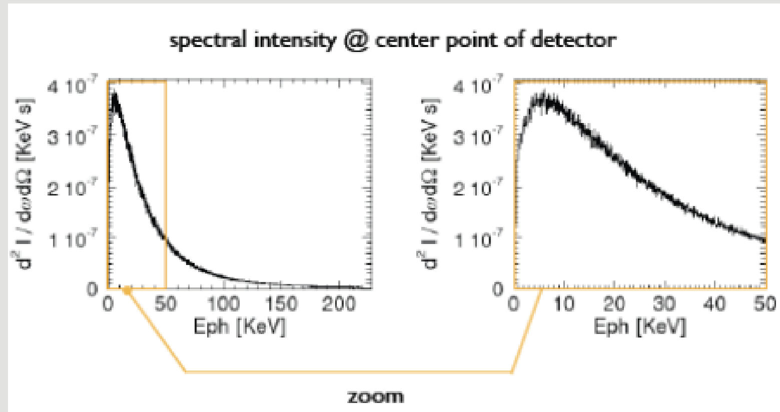

Successful benchmarking of 3D LWFA simulations on ~1,000 procs

Time-explicit PIC (OSIRIS, VORPAL), quasi-static (QuickPIC) and laser envelope (VORPAL) results agree for an intense laser pulse entering a uniform plasma –

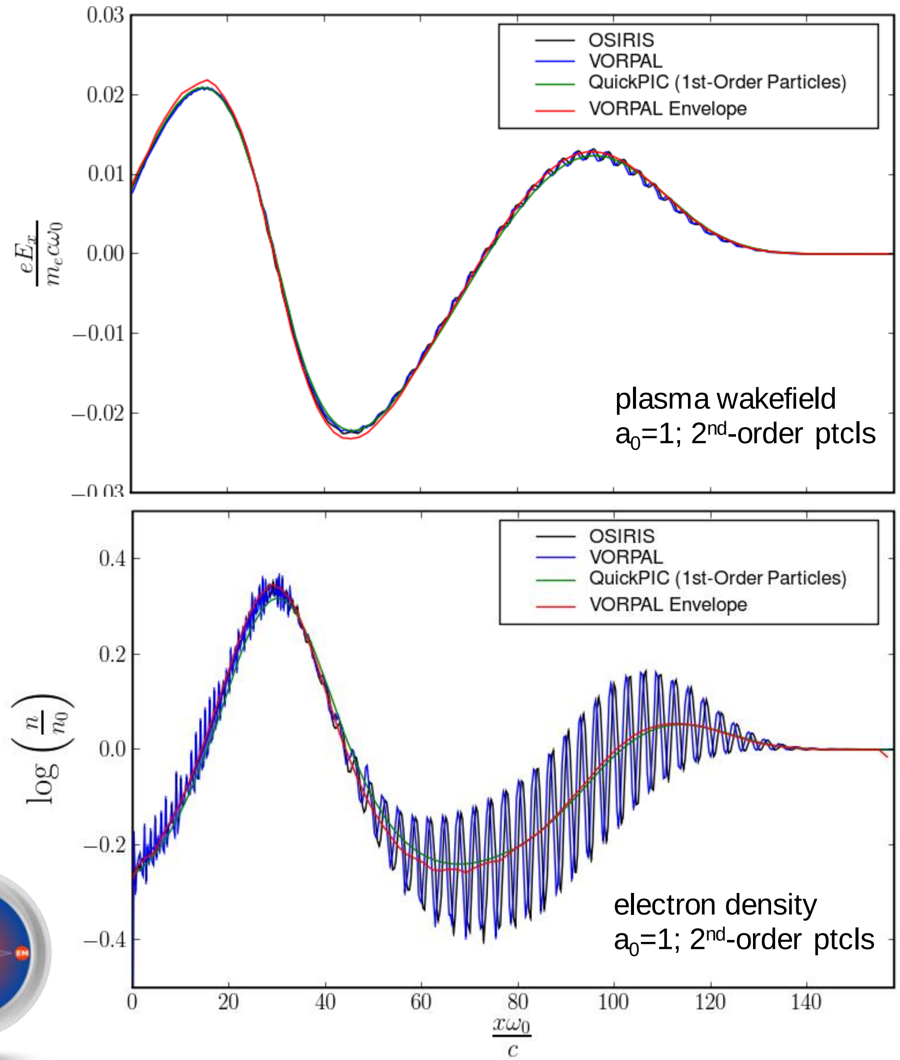

Physical parameters:

τ_{fwhm} = 30 fs
w_0 = 8.2 μm
P_0 = 2.26 TW
I_{peak} = 2.14×10^{18} W cm^{-2}
a_0 = 1.0
n_e = 1.38×10^{19} cm^{-3}

Explicit numerical parameters:

N_{cells} = 512 x 512 x 512
= 1.34×10^8
N_{ptcls} = 1.07×10^9
Δ_{trans} = 0.159 μm
Δ_{long} = 0.04 μm = $\lambda_0/20$
δt = 0.0998 fs

TECH-X CORPORATION

Successful benchmarking of 3D LWFA simulations on ~1,000 procs

OSIRIS – time-explicit PIC

QuickPIC – quasi-static PIC

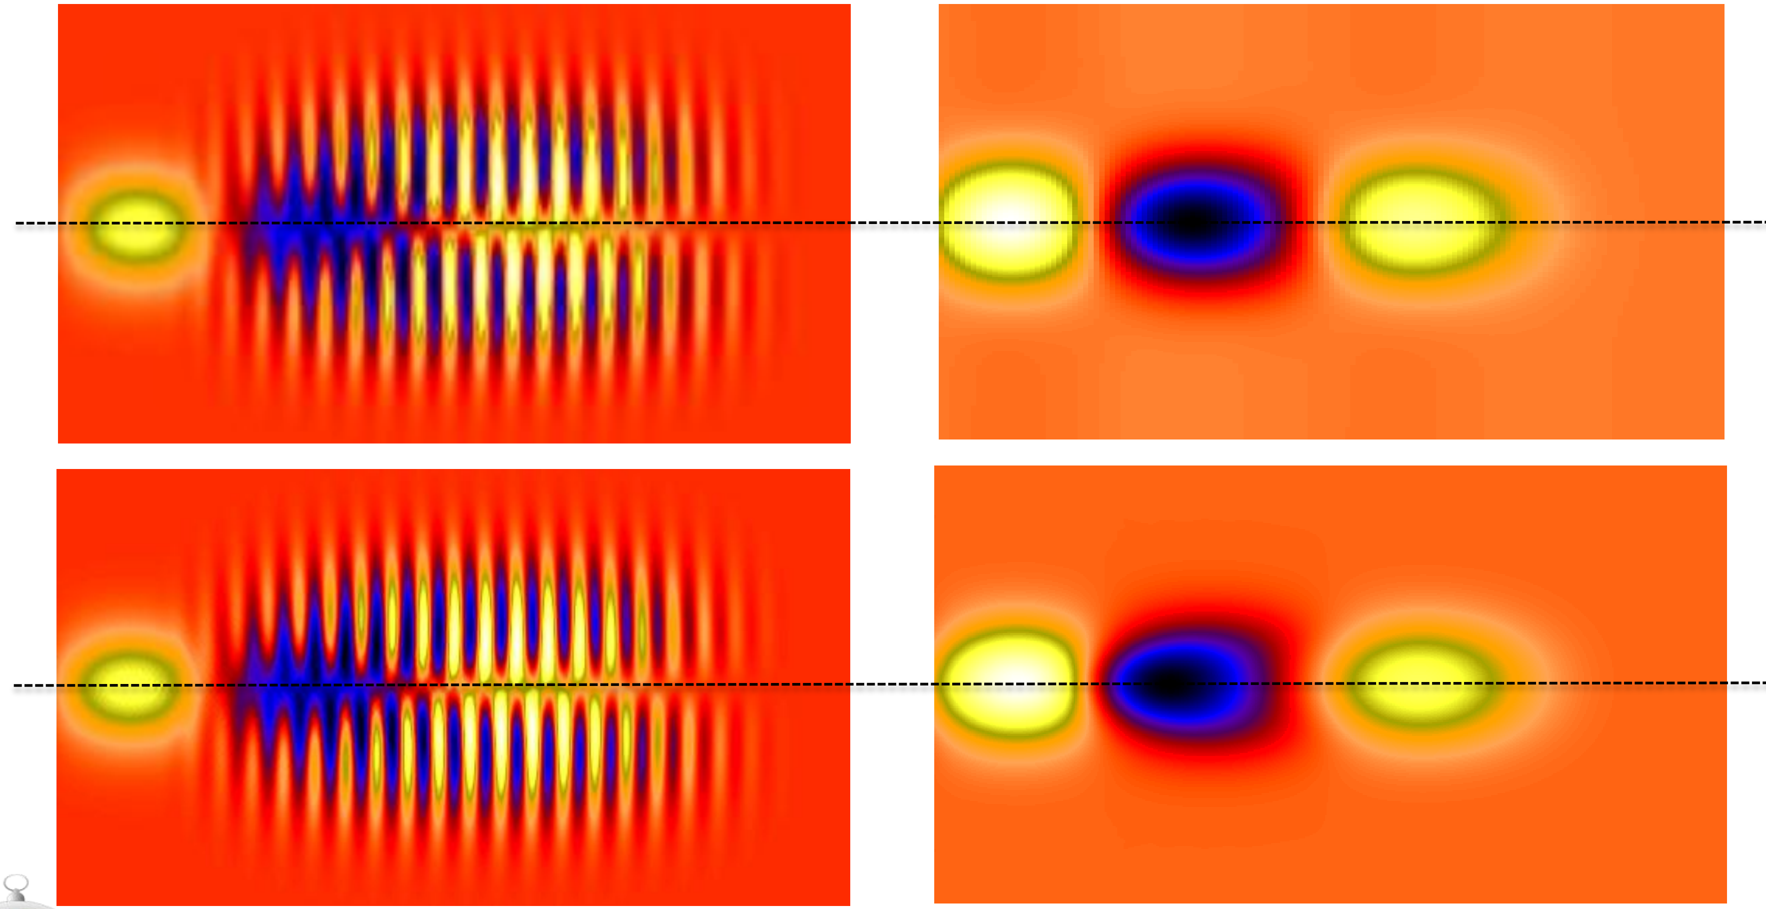

VORPAL – time-explicit PIC

VORPAL – laser envelope model

TECH-X CORPORATION

New schemes and structures required

New sorted PIC algorithm

Advantages of this new structure

- Single precision can be used for particles
- Reduced memory bandwidth requirements
- Reduced need for cache
- Elimination of indirect memory addressing (gather/scatter) for fields
- Allows for fine grained partitioning and load balancing

Disadvantages of this new structure

- Maintaining particle order adds complexity

Issues with single precision

Tap into the power of state of the art processing units

- Generally limited to single precision arithmetic
- Specific C/C++ code

Verify impact of numerical precision

- Fields, Particles
- Positions defined as cell index + position within cell

Interface with hardware specific code

- Compatible with existing code structure
- Write hardware optimized routines

GPU & SIMD optimizations and benchmarks

UCLA

GPUs

GPUs are graphical processing units which consist of

- 12-30 multiprocessors, each with a small (16KB), fast (4 clocks) shared memory
- Each multi-processor contains 8 processor cores
- Large (0.5-4.0 GB), slow (400-600 clocks) global memory, readable by all units
- No cache
- Very fast (1 clock) hardware thread switching

2D Charge Deposit Benchmark

64x128 grid, 294,912 particles, 36 particles/cell

Results compared to the 3.0 GHz Intel Host

- GTX 280 gave a **speedup of 25**

3D PIC loop

Estimated speedup of **~100**

SIMDs & OSIRIS

Single Instruction Multiple Data

- Modern cpus (Intel/AMD/PowerPC) include a SIMD vector unit
- Vector registers (4x 32 bit int/float)*
- Instructions act on vector registers (4 simultaneous operations)
- Require ASM or C intrinsics
- Usable on most C compilers (gcc, icc, etc.)
- Same concepts apply to PowerPC Altivec units

2D PIC Loop Benchmark for OSIRIS

Normalized performance **for quadratic splines**

128x128 grid, ~8M particles, 512 particles/cell

Results compared on a 3.2 GHz Intel i7

- **F90** gives a PIC loop of **155 ns**/part/step
- **SSE** gives a PIC loop of **89 ns**/part/step

Planned VORPAL enhancements: better messaging for strong scaling; optimize particle push for single processor performance; port fields/particles to GPU

VORPAL enhancements on Petascale systems: strong scaling and single processor performance

Peter Messmer,[1] Ben Cowan,[1] George Bell,[1] Keegan Amyx,[1] Boyana Norris[2] & John R. Cary[1]

[1]*Tech-X Corp.*, [2]*Argonne National Lab.*

Supported by DOE/ASCR SBIR: DE-FG02-07ER84731 & VORPAL customers

- Work on field messaging enables 10x10x10 domain sizes (see J. Cary presentation)
- Development and implementation of optimized particle push is in progress
  - 0.12 μs/ptcl/step (2.3 GHz opteron) is achieved (C/MPI test kernel; no deposition)
    - explicit vectorization, optimization of data layout, tuning compiler optimization flags
  - 0.2 μs/ptcl/step (2.3 GHz opteron) is the goal (VORPAL, w/ current deposition, double precision)
  - 0.08 μs/ptcl/step has been achieved in VPIC (LANL, single-precision with altivec instruction set)

NVIDIA GPU acceleration of FDTD simulations with conformal boundaries

Peter Messmer,[1] Travis Austin,[1] John R. Cary,[1] Paul Mullowney,[1] Keegan Amyx[1] & Mike Galloy[1]

[1]*Tech-X Corp.*

Partially supported by NASA SBIR # NNG06CA13C, NVIDIA Corp. & Tech-X Corp.

- 3D electromagnetics with conformal boundaries & dielectrics has been implemented
  - available in high-level languages (Matlab, IDL, python), as well as C/C++
  - accomplished via GPULib — **http://gpulib.txcorp.com/** — 20x speedup observed
    - 3D domain unwrapped into 1D vector; extra layer of guard cells; BC cleanup via "dielectric mask"
- Particle push without current deposition has been prototyped
  - potential race conditions have been identified for current deposition
    - ideas to solve these problems are waiting to be tested (no funding at present)
- Implementation in VORPAL is resource limited
  - needed for future, heterogeneous architectures

TECH-X CORPORATION

Future Plans

- Continue experimental support of LBL / BELLA and physics discovery
  - physics of self-trapping and controlled injection
  - 10 GeV stages for e- and e+ acceleration, with emittance control
  - also, more nonlinear regimes, validate against other experiments
  - Provide high-fidelity modeling for planning and optimizing BELLA experiments
- Develop comprehensive LWFA simulation capability and explore collider options
  - model high visibility experiments
  - meter-scale plasmas,   e- and e+ acceleration
  - controlled optical injection of e- beams.
  - compare beam loading schemes
  - accurate evolution of low-emittance beams
  - nonlinear vs. weakly nonlinear vs. quasi-linear regimes
  - Staging, guiding, pulse shaping, stability control
- Continue code verification and validation
- Continue enhance suite of approaches
  - Improve speed up of quasi-static, noise reduction, improved dispersion, mesh refinement, radiation models, reduced models, high-order FDTD, cold relativistic fluid…
- Continue VACET collaboration on rapid 3D viz and post-processing
- Continue to improve parallel scaling and code efficiencies
  - Develop PIC algorithms for advanced architectures